

# Red Hat OpenShift Container Storage

Persistent and highly available storage for Red Hat OpenShift logging

## Highlights

Deploy persistent software-defined Red Hat OpenShift

Container Storage for Red Hat OpenShift Container Platform logging.

Replicate log data across multiple Amazon Web Services (AWS) Availability Zones for failover and resilience.

Maintain performance and ongoing operations for Red Hat OpenShift Container Platform services and key applications, even across storage failures.

Serve Red Hat OpenShift Container Platform applications with a choice of block, file, or object storage.

## Persistent and highly available storage

As organizations embrace containers and Kubernetes orchestration for critical applications, Red Hat® OpenShift® Container Platform services must be made highly available, implying a significant role for the storage that supports them. Red Hat OpenShift Container Platform cluster logging requires persistent storage for reliable operation. Back-end storage for logging services must also be resilient and highly available to protect against infrastructure failures without data loss.

Red Hat OpenShift Container Storage is an ideal software-defined storage back end for Red Hat OpenShift logging because it offers data persistence and storage resiliency across failure domains such as Amazon Web Services (AWS) Availability Zones or on-prem systems. OpenShift Container Storage provides a uniform experience across on-prem, public cloud, or hybrid cloud deployments while offering block, file, and object storage classes to support a wide range of applications. As described in this document, testing has verified that OpenShift Container Storage provides superior failover protection for Red Hat OpenShift logging with no performance bottleneck.

## Red Hat OpenShift Container Storage for Red Hat OpenShift logging

Red Hat OpenShift Container Platform administrators can select from multiple back-end storage options for Red Hat OpenShift logging. By default, OpenShift Container Platform uses ephemeral [emptyDir volumes](#) for log storage. When a pod is removed from a node for any reason, the related non-persistent data in the emptyDir is removed as well. AWS Elastic Block Storage (EBS) persistent volumes (PVs) provide an alternative to emptyDir in the Amazon public cloud. Although they offer failover for log storage within a single availability zone, EBS PVs provide no support for data storage or failover across multiple availability zones. OpenShift Container Storage addresses both of these issues.

## Red Hat OpenShift Container Platform logging

Administrators can deploy cluster logging using the OpenShift Container Platform web console to install the [Elasticsearch operator](#) and the [cluster logging operator](#). The cluster logging components are based on the Elasticsearch, Fluentd, and Kibana (EFK) stack:

- ▶ **Elasticsearch index** distributes a collection of documents across different containers. Pieces of Elasticsearch indexes (shards) are duplicated across a set of nodes to provide redundant copies (replicas) in case of infrastructure failure.
- ▶ **Fluentd** acts as a collector when deployed to each node in the Red Hat OpenShift Container Platform cluster. It collects node and container logs and writes them to Elasticsearch.
- ▶ **Kibana** is a centralized web user interface where users and administrators can create rich visualizations and dashboards from aggregated data.



[facebook.com/redhatinc](http://facebook.com/redhatinc)  
@RedHat  
[linkedin.com/company/red-hat](http://linkedin.com/company/red-hat)

Designed for and with  
 Red Hat OpenShift  
 Container Platform,  
 Red Hat OpenShift  
 Container Storage  
 offers highly  
 scalable, persistent  
 storage for cloud-  
 native applications.

### Comparing backend storage options

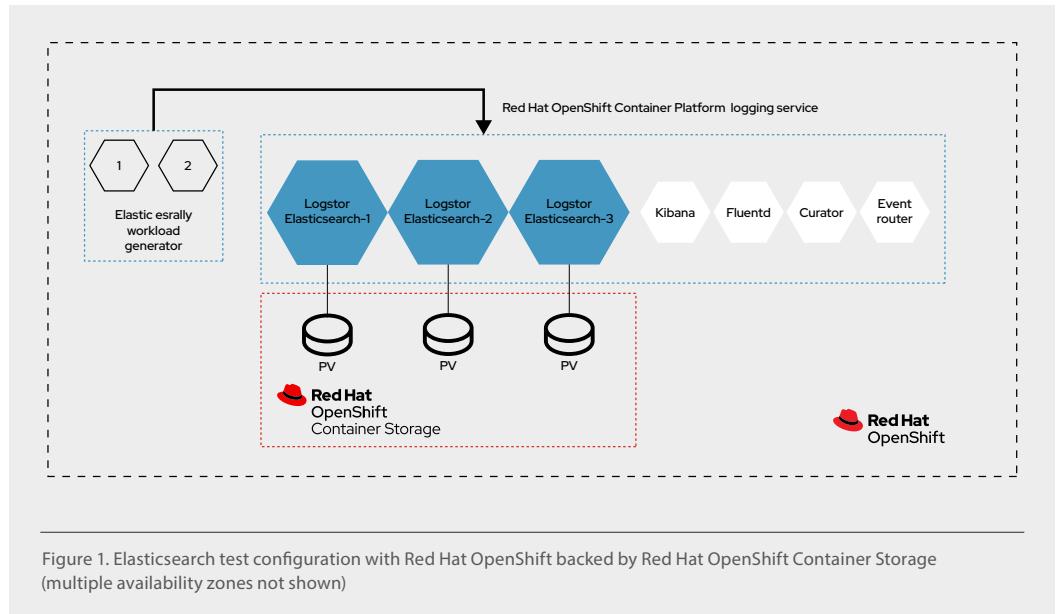
Administrators and developers should consider multiple factors when deploying either Red Hat OpenShift cluster logging or general-purpose Elasticsearch on Red Hat OpenShift Container Platform. Elasticsearch offers optional replication at the index level to provide data resilience. Additional data resilience can be provided by deploying Elasticsearch on top of a reliable storage service such as OpenShift Container Storage. Based on Red Hat testing, Table 1 compares OpenShift Container storage with AWS EBS both with and without Elasticsearch redundancy. Scores are ranked based on the following values from low to high performance: fair (1), good (2), better (3), best (4).

**Table 1. Red Hat OpenShift cluster logging feature comparison (Elasticsearch)**

|                               | Zero Elasticsearch redundancy (default) |        | Single Elasticsearch redundancy |               |
|-------------------------------|---|--------|---------------------------------|---------------|
|                               | OpenShift Container Storage             | EBS    | OpenShift Container Storage     | EBS           |
| Use case                      | Logs, metrics, APM, log analytics       |        | General purpose Elasticsearch   |               |
| Elasticsearch resiliency      | Best                                    | Fair   | Best                            | Fair          |
| Single availability zone HA   | Best                                    | Best   | Best                            | Best          |
| Multiple availability zone HA | Best                                    | Fair   | Best                            | Not available |
| Performance                   | Better                                  | Better | Good                            | Good          |
| Reliability                   | Best                                    | Better | Best                            | Best          |

### Workload testing summary

The test environment consisted of a Red Hat OpenShift Container Platform cluster configured with both OpenShift Container Storage and AWS EBS storage classes configured across three AWS Availability Zones (shown in Figure 1 with cluster details listed in Table 2).

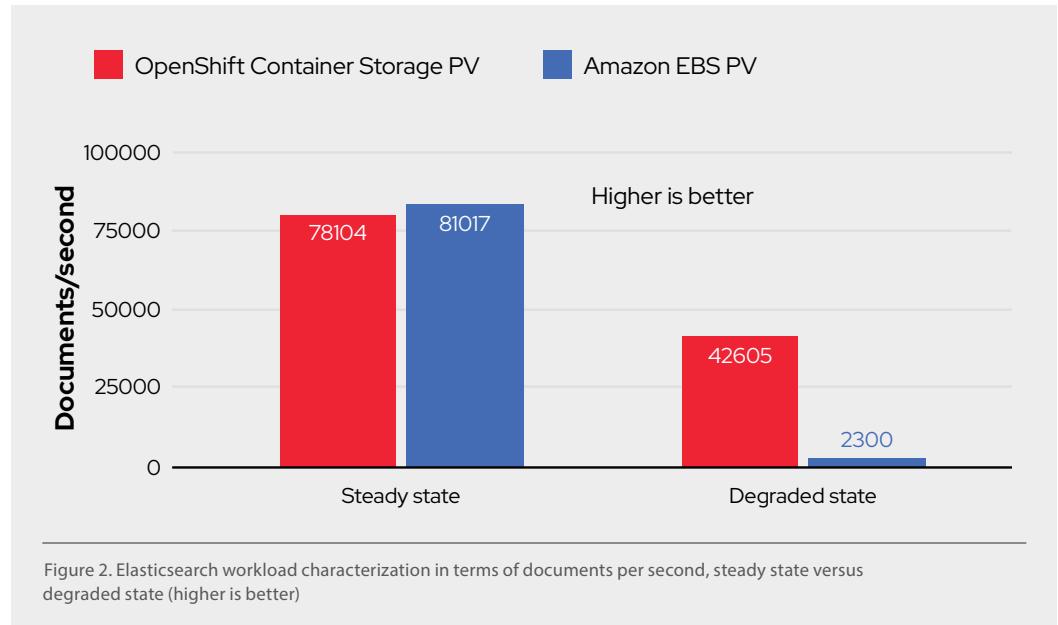


**Table 2. Test cluster details<sup>1</sup>**

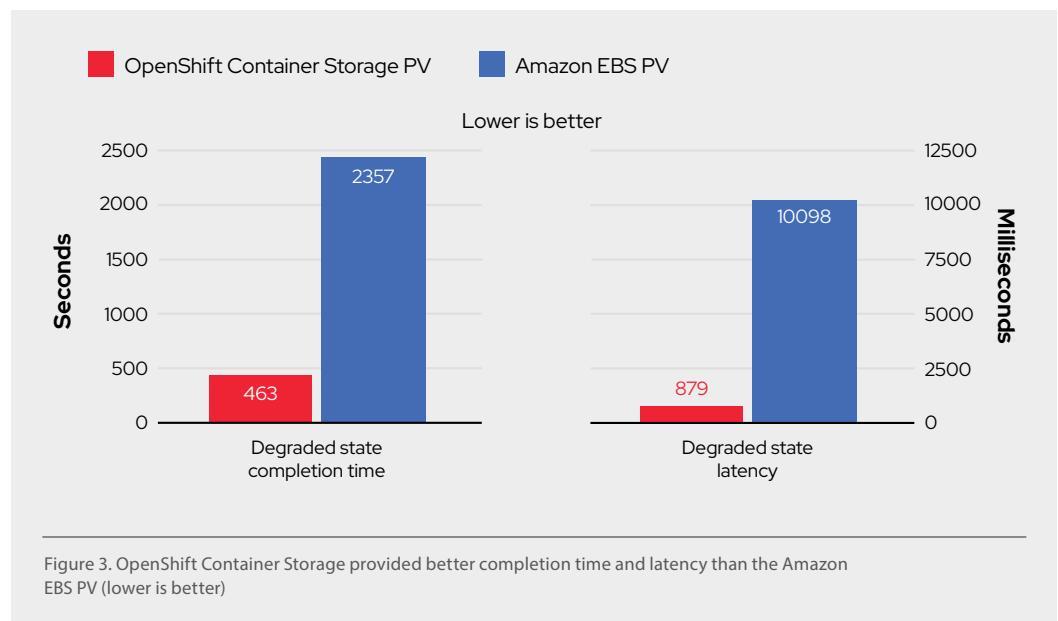
|                                      |  |
|--------------------------------------|--|
| Master nodes                         | 3x m4.large instances                                    |
| Elasticsearch host nodes             | 4x AWS m4.2xlarge instances<br>(1 per Availability Zone) |
| OpenShift Container Storage nodes    | 3x i3en.2xlarge (1 per Availability Zone)                |
| OpenShift Container Storage devices  | 1x 1.9TB NVMe / node                                     |
| Persistent volumes (PVs)             | 3x 200GB PVs (1 per Elasticsearch pod)                   |
| Red Hat OpenShift Container Platform | Version 4.3  |
| Cluster logging Operator version     | 4.3.7  |
| Elasticsearch version                | 5.6.16   |

To measure performance and resilience for the logging service, Red Hat engineers simulated Elasticsearch node failure while a client was performing an index append operation to the cluster. As shown in Figure 2, steady state performance was roughly equivalent between Elasticsearch PVs running on OpenShift Container Storage and AWS EBS. Under node failure conditions (degraded state), the cluster running on OpenShift Container Storage continued to operate normally, delivering consistent index append performance (higher is better). The index-append error rate was zero for the OpenShift Container Storage cluster, showing full redundancy and uninterrupted service.

<sup>1</sup> For details, refer to the blog [OpenShift Container Storage 4 and logging with Elasticsearch and cluster logging operators](#).



In the degraded state, Elasticsearch deployed on the AWS EBS storage class generated an index append error rate of 100% and lower index-append throughput, along with longer completion time and higher latency than the OpenShift Container Storage PV (Figure 3).



This degraded resiliency during infrastructure failure (in this case, node failure) is explained by EBS storage class design. EBS is a zonal service and is limited to a single AWS Availability Zone within a region. As such, when Red Hat OpenShift Container Platform spawns a new Elasticsearch pod in a different availability zone, the new pod is unable to access the existing persistent volume. The Elasticsearch cluster enters an error (ERR) state as a result.

When Red Hat OpenShift Container Platform is deployed in the public cloud, Red Hat OpenShift nodes can be deployed in multiple availability zones within a region, forming a distributed and highly available storage cluster. Storage provisioned by OpenShift Container Storage is also available across availability zones. As such, OpenShift Container Storage can handle the failure of an entire availability zone without losing data or impacting application service availability – ultimately yielding higher resilience for applications consuming Red Hat OpenShift Container Storage services.

Persistent and highly available storage across multiple availability zones for Red Hat OpenShift Container Platform logging means Elasticsearch can withstand a variety of infrastructure failure scenarios, including:

- ▶ Failure in an Elasticsearch pod
- ▶ Failure in an Elasticsearch pod host node
- ▶ Failure in an Elasticsearch storage PV
- ▶ Failure of the common storage backend
- ▶ Failure of the infrastructure availability zone or site

## Conclusion

Providing high availability for container-based applications requires a Red Hat OpenShift Container Platform environment that is, itself, highly available, which requires highly available persistent storage. Red Hat OpenShift Container Storage supports Red Hat OpenShift Container Platform logging services by providing software-defined persistent storage as well as critical redundancy across AWS Availability Zones. The result is logging services that can continue to function despite infrastructure failures. The same OpenShift Container Storage deployment can also support Red Hat OpenShift applications with container-native file, block, or object storage.



## About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.



[@RedHat](https://facebook.com/redhatinc)  
[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

[#F25144\\_0920](https://redhat.com)

**North America**  
1888 REDHAT1  
[www.redhat.com](http://www.redhat.com)

**Europe, Middle East, and Africa**  
00800 7334 2835  
[europe@redhat.com](mailto:europe@redhat.com)

**Asia Pacific**  
+65 6490 4200  
[apac@redhat.com](mailto:apac@redhat.com)

**Latin America**  
+54 11 4329 7300  
[info-latam@redhat.com](mailto:info-latam@redhat.com)